
Robust Classification with Adiabatic Quantum Optimization

Vasil S. Denchev

Purdue University/Google

DENCHEV@GMAIL.COM

Nan Ding

Purdue University

DING10@PURDUE.EDU

S. V. N. Vishwanathan

Purdue University

VISHY@STAT.PURDUE.EDU

Hartmut Neven

Google

NEVEN@GOOGLE.COM

Abstract

We propose a non-convex training objective for robust binary classification of data sets in which label noise is present. The design is guided by the intention of solving the resulting problem by adiabatic quantum optimization. Two requirements are imposed by the engineering constraints of existing quantum hardware: training problems are formulated as quadratic unconstrained binary optimization; and model parameters are represented as binary expansions of low bit-depth. In the present work we validate this approach by using a heuristic classical solver as a stand-in for quantum hardware. Testing on several popular data sets and comparing with a number of existing losses we find substantial advantages in robustness as measured by test error under increasing label noise. Robustness is enabled by the non-convexity of our hardware-compatible loss function, which we name q -loss.

to that. While it is easily recognized that allowing for non-convex objectives opens up a plethora of possibilities for better solutions to machine learning problems, much of the contemporary research has deliberately avoided them. The reason is the widely known fact that non-convexity often results in NP-hard problems.

However this choice comes at a cost as the shortcomings of convex objectives are also well understood. Recent work (Long & Servedio, 2010) showed that convex loss functions cannot be made robust in the presence of label noise because they cause unbounded growth of penalties for large negative margins. (Manwani & Sastri, 2011) further characterized this effect by analyzing various convex losses and found that none of them is tolerant to non-uniform label noise. In practice label noise turns out to be a serious problem due to the fact that it affects real-world data sets to a significant degree. Since label noise manifests itself throughout the optimization as large negative margins, the finally constructed decision hyperplane that represents the global minimum of any convex loss tends to be pulled by the mislabeled training examples away from the minimizer of classification error. Therefore, even though solving convex losses to optimality is feasible, when label noise causes the lowest objective value to not correspond to the lowest attainable training error, the entire exercise misses the mark. Consequently, any approach exhibiting this problem does not stand to benefit from improved optimization techniques.

For example, Fig. 1 shows the broken correspondence between training error and objective value when a convex loss is used in a training problem of practical significance—"OCR in photos". The human task of tagging characters in photos of potentially poor qual-

1. Introduction

In recent years machine learning researchers and practitioners have been focusing on convex optimization methods due to their computational advantages and well understood mathematical properties. The many successes of convexity-based algorithms are witnesses

Appearing in *Proceedings of the 29th International Conference on Machine Learning*, Edinburgh, Scotland, UK, 2012. Copyright 2012 by the author(s)/owner(s).




Figure 1. Relationship between training error and inverse empirical risk produced by minimizing square loss on six different binary classifiers for digits (e.g. '1' vs the rest, '2' vs the rest, etc.) The data ("OCR in photos"; 10200 dimensions; 38924 examples; 10 classes) represents a challenging real-world training problem of significant practical importance. An adequate loss function should generally be decreasing training error as the empirical risk approaches global minimum (top plots). Unfortunately, the opposite effect (bottom plots) can often be observed when working with convex losses. The failures are found to be due to two factors, both of which cause square loss to be drastically misled by its convexity: occasionally mistaken labels resulting from the semi-automatic process generating the data; and the presence of examples of one class that may be similar to examples of another class (e.g. '6' and '8').

ity is not easy, so the presence of mislabeled examples in the training set is not surprising. Even worse, routinely used semi-automatic preparation of training data is also contributing to mistakes. The problem may gradually disappear for cleaner data sets, which often happen to be the cases when convex losses produce excellent classifiers. Unfortunately the nature of large-scale supervised learning does not permit elaborate quality assurance for data sets that are handed out to training algorithms; accordingly label noise will continue to pollute real-world data sets. Moreover future intelligent systems will rely increasingly on weakly labeled or unlabeled data increasing the need for noise tolerance.

(Ding & Vishwanathan, 2010) and (Masnadi-Shirazi et al., 2010) took these lessons and independently studied two different non-convex but seemingly well behaved types of loss functions. (Collobert et al., 2006; Ertekin et al., 2011) also explored non-convexity in the context of SVM with ramp loss, but their focus was on achieving sparser sets of support vectors and speed of training rather than improved accuracy and robustness of the constructed classifier.

In the present work we continue the study of non-convexity. We report on training with a non-convex objective using discrete optimization in a formula-

tion adapted to take advantage of emerging hardware that performs adiabatic quantum optimization (AQO). AQO, first proposed in (Farhi et al., 2000), is a quantum computing model with good prospects for scalable and practically useful hardware implementation. Studies of its purported computational superiority over classical computing have repeatedly given encouraging results, e.g. (Dickson & Amin, 2011). Significant investments are underway by the Canadian company D-Wave to develop a hardware implementation. A series of rigorous studies of the quantum mechanical properties of the D-Wave processors, culminating in a recent Nature publication (Johnson et al., 2011), have increased the excitement in the quantum computing community for this approach. This was further fueled by news of a successful collaboration with Google (Neven et al., 2009a) and of Lockheed Martin purchasing an adiabatic quantum computer. For machine learning purposes, D-Wave's implementation of AQO can be regarded as a black-box discrete optimization engine that accepts any problems formulated as quadratic unconstrained binary optimization (QUBO), also equivalent to the Ising model and Weighted MAX-2-SAT. It should be noted that this training formulation is a good format for AQO independently of D-Wave's efforts since it can be physically realized as the simplest possible multi-qubit configuration—an Ising system (Brush, 1967). We do not claim principled superiority of q -loss over other non-convex losses. However q -loss is distinguished by the fact that it can be formulated for AQO on quantum hardware that only supports quadratic (2-local) interactions among its qubits using a number of ancillary variables that just grows linearly with the number of training examples. To the best of our current knowledge, no other non-convex loss has this property¹. While all other non-convex losses are tackled by heuristic optimization with very limited success, q -loss may be solvable to optimality by AQO.

The paper is organized as follows: Section 2 defines the training problem; Section 3 introduces q -loss, derives its QUBO formulation, and discusses the intuition behind it; Sections 4 and 5 deal with choosing hyper-parameter values and discretization of variables; Section 6 presents our experiments; and Section 7 concludes with an overview and discussion. Technical details can be found in the supplementary material.

¹Except for the non-margin-enforcing 0-1 loss (Neven et al., 2009b)

2. Training a binary classifier

We study binary classifiers $y = \text{sign}(\mathbf{w}^T \mathbf{x} + b)$, where $\mathbf{x} \in \mathbb{R}^N$ is an input pattern to be classified, $y \in \{-1, 1\}$ is the label associated with \mathbf{x} , $\mathbf{w} \in \mathbb{R}^N$ is a vector of weights to be optimized, and $b \in \mathbb{R}$ is the bias. Training, also known as regularized risk minimization, consists of choosing \mathbf{w} and b by simultaneously minimizing two terms: *empirical risk* $R(\mathbf{w}, b) = \sum_{s=1}^S L(m(\mathbf{x}_s, y_s, \mathbf{w}, b)) / S$ and *regularization* $\Omega(\mathbf{w})$. R , via a *loss function* L , estimates the error that any candidate classifier causes over a set of S training examples $\{(\mathbf{x}_s, y_s) | s = 1, \dots, S\}$. The argument of L is known as the *margin* of example s with respect to the decision hyperplane defined by \mathbf{w} and b :

$$m(\mathbf{x}_s, y_s, \mathbf{w}, b) = y_s(\mathbf{w}^T \mathbf{x}_s + b) \quad (1)$$

Ω controls the complexity of the classifier and is necessary for good generalization because classifiers with high complexity display overfitting—they can classify the training set with low error but may not do well on previously unseen data. Training amounts to solving

$$(\mathbf{w}, b)^* = \arg \min_{\mathbf{w}, b} \{R(\mathbf{w}, b) + \Omega(\mathbf{w})\} . \quad (2)$$

The most natural choice for L is *0-1 loss*, which simply indicates a misclassification for a negative margin:

$$L_{0-1}(m) = (1 - \text{sign}(m)) / 2 \quad (3)$$

Due to the non-convexity of L_{0-1} , the resulting optimization problem (2) is NP-hard (Feldman et al., 2010). To avoid dealing with NP-hard optimization problems, in practice L_{0-1} is replaced by some convex upper bound (e.g. square, logistic, exponential, hinge), and Ω is usually chosen as ℓ_1 - or ℓ_2 -norm penalization of \mathbf{w} . This allows arriving at convex optimization problems that can be rigorously analyzed and efficiently solved by classical means. However, such relaxations are known to compromise the original goal of training because convex losses can be severely misled by label noise in the training data.

3. q -loss

Because the quantum hardware natively represents a general family of quadratic functions, the simplest loss function that would work is *square loss*, which is a convex upper bound to L_{0-1} :

$$L_{\text{square}}(m) = (m - 1)^2 \quad (4)$$

However, there are two drawbacks of square loss when applied to binary classification. First, in binary classification it does not make sense to penalize large positive margins. Second, as mentioned earlier, square

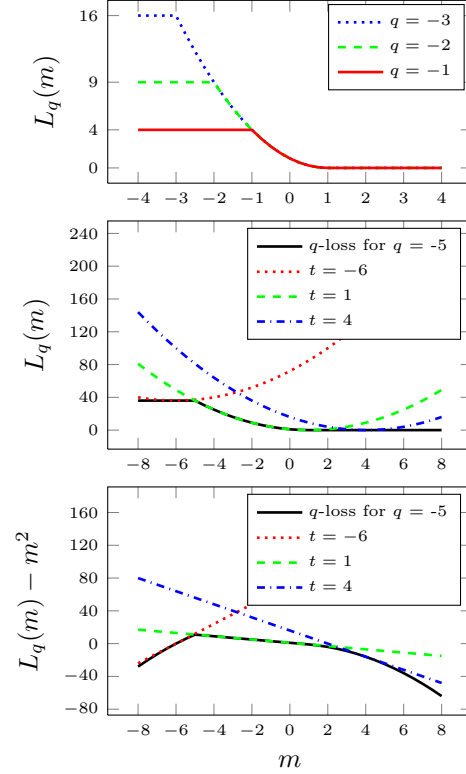


Figure 2. Top: q -loss for different values of q . Middle: q -loss with three members of the quadratic upper bounds family. $t \in \mathbb{R}$ is the variational parameter. Bottom: Transforming the y -axis for concavity.

loss has the same flaw as all convex losses—penalties for large negative margins grow unboundedly, which can cause non-robustness with respect to label noise.

With these considerations in mind, we modify square loss in order to obtain a training formulation for binary classification that is both compatible with quantum hardware and robust to label noise. The resulting loss, which we name *q -loss* (Fig. 2, top), is essentially a doubly truncated version of (4) with parameterization over $q \in (-\infty, 0]$ defined as follows:

Definition 1 (q -loss)

$$L_q(m) = \min \left((1 - q)^2, (\max(0, 1 - m))^2 \right) \quad (5)$$

Unfortunately, (5) does not lead to a QUBO. However, it turns out that we can transform it into a problem which can be solved as a QUBO. The basic idea is to find a variational approximation via a family of quadratic functions that upper-bound q -loss and are governed by a variational parameter $t \in \mathbb{R}$ as shown in Fig. 2, middle.

Theorem 2 q -loss in (5) is equivalent to:

$$L_q(m) = \min_t \left\{ (m-t)^2 + (1-q)^2 \frac{(1 - \text{sign}(t-1))}{2} \right\}$$

Proof Since q -loss is non-convex, the standard derivation via convex duality (Jordan et al., 1999) dictates that we first find a new coordinate system in which q -loss is concave or convex. Then we calculate the conjugate function for linear bounds in the transformed space and transform back to the original space where the linear bounds become the quadratic bounds shown in Fig. 2, middle. Because of the presence of two constant segments in q -loss, any coordinate system in which the two axes are independent transformations of the original x and y axes clearly cannot result in concavity or convexity. Thereby we are led to the transformation $f(y) = y - x^2$, which gives $f(L_q(m)) = L_q(m) - m^2$. It can be seen (Fig. 2, bottom) that in this transformed space q -loss is concave and the quadratic upper bounds become tangent lines. The conjugate function in the transformed space is $g(\eta) = \min_m \{\eta m - f(L_q(m))\}$.

To minimize, we seek stationary points by differentiating $\phi(\eta, m) = \eta m - f(L_q(m))$ with respect to m :

$$\begin{aligned} \frac{\partial}{\partial m} \phi(\eta, m) &= \eta - \frac{d}{dm} L_q(m) + 2m \\ &= \begin{cases} \eta + 2 & \text{for } m \in (q, 1) \\ \eta + 2m & \text{for } m \in (-\infty, q) \cup (1, \infty) \end{cases}, \end{aligned} \quad (6)$$

as yielded by piecewise differentiation of $L_q(m)$. Setting to 0 gives the stationary points

$$\eta = -2 \text{ for } m \in (q, 1) \quad (7)$$

$$m = -\eta/2 \text{ for } m \in (-\infty, q) \cup (1, \infty). \quad (8)$$

Plugging them back into the conjugate function yields

$$\begin{aligned} g(\eta) &= \begin{cases} -\frac{\eta^2}{4} - (1-q)^2 & \text{for } m \in (-\infty, q) \\ -1 & \text{for } m \in (q, 1) \\ -\frac{\eta^2}{4} & \text{for } m \in (1, \infty) \end{cases} \\ &= -\frac{\eta^2}{4} - (1-q)^2 \frac{(1 - \text{sign}(-\frac{\eta}{2} - 1))}{2}. \end{aligned} \quad (9)$$

In accordance with convex duality,

$$\begin{aligned} f(L_q(m)) &= \min_{\eta} \{\eta m - g(\eta)\} \\ &= \min_{\eta} \left\{ \eta m + \frac{\eta^2}{4} + (1-q)^2 \frac{(1 - \text{sign}(-\frac{\eta}{2} - 1))}{2} \right\}. \end{aligned} \quad (10)$$

Transforming back into the original space and setting $t = -\eta/2$, the variational upper bound for q -loss is

$$\begin{aligned} L_q(m) &= f^{-1}(f(L_q(m))) \\ &= \min_t \left\{ (m-t)^2 + (1-q)^2 \frac{(1 - \text{sign}(t-1))}{2} \right\} \end{aligned} \quad (11)$$

■

3.1. Latent variables view

Traditionally when facing non-convex optimization problems, a viable approach is to introduce latent variables that allow reformulating over a simpler family of functions. This is precisely what Theorem 2 achieves. For any fixed m , the latent variable $t \in \mathbb{R}$ gives a convex optimization problem whose minimum is $L_q(m)$:

$$L_q(m) = h(m, t^*(m)), \text{ where} \quad (12)$$

$$t^*(m) = \arg \min_t \{h(m, t)\}$$

$$h(m, t) = (m-t)^2 + (1-q)^2 (1 - \text{sign}(t-1)) / 2$$

The regularized risk minimization (2) with empirical risk over L_q in the form (12) is amenable to a block coordinate descent method for jointly optimizing the model parameters (\mathbf{w}, b) and the latent variables t_s for $s = 1, \dots, S$: similarly to EM, alternate between convex optimization runs over the latent variables (t step) and the model parameters (w step). Even though such methods do well on some problems with certain benign structure—e.g. Gaussian mixtures (Dempster et al., 1977)—they are also known to fail on other problems that lack such structure. We believe q -loss belongs to the latter group and have verified that a block coordinate descent method is likely to be sensitive to initialization and is quickly terminating in bad local minima. The intuitive reason is that due to the quadratically growing penalty for mismatching a margin with its latent variable, the t step tends to lock in the model parameters found during the previous w step, thus possibly preventing the next w step from moving to a different model. The impact of this effect becomes ever more severe for large data with $S \gg N$.

On the other hand, by transforming (5) into (12) we have made training with q -loss representable in QUBO form albeit at the expense of additional variables. Section A of the supplementary material explicitly shows the QUBO problem that can be derived from (12). Since the goal of AQO is to perform global optimization simultaneously over all variables, we believe AQO is a much better candidate for training with q -loss. Besides making the QUBO formulation possible, the

introduction of latent variables also gives rise to an intuitive interpretation of the mechanism by which q -loss achieves robustness when compared to the non-robustness of square loss. While in (4) the fixed target 1 has to be matched as closely as possible by m , in (12) t plays the role of a flexible target that can change sign for a large negative margin, thereby flagging that training example as mislabeled. For any m , the minimizer $t^*(m)$ in (12) belongs to one of three cases:

- Case I: $m \geq 1 \Rightarrow t^*(m) = m$
- Case II: $q < m < 1 \Rightarrow t^*(m) = 1$
- Case III: $m \leq q \Rightarrow t^*(m) = m$

Case I ensures zero penalty for large positive margins; Case II produces the same quadratic penalty as (4); Case III can be seen as flipping the label of a possibly mislabeled example but also incurring a constant penalty of $(1-q)^2$ in order to not lose connection with the original labeling. Thus, the hyper-parameter q defines the largest negative margin to be tolerated. A training example that has a negative margin with some larger magnitude gets flipped with constant penalty.

4. Bounding q

While it is difficult to formalize any general statements about the computational hardness of q -loss, it is easily recognized that the hardness depends on the size of the parabolic segment controlled by q . For $q \rightarrow -\infty$, even the negative margins of highest magnitude incur the usual quadratic penalty, and the loss becomes effectively convex. For smaller q the loss becomes similar to 0-1 loss, so the resulting optimization problems may be approaching the hardness of the corresponding 0-1 loss problems. However, the most beneficial regime of operation is not known a-priori. This necessitates cross-validation over q , which, depending on the noise level, we expect to result in some trade-off between robustness and computational hardness. For the purpose of choosing values for cross-validation, we give an approximate lower bound for q as a function of our estimate of the underlying Bayes error in the data and the label noise that we might artificially insert into the training set for robustness evaluation.

Let the effective Bayes error be $\beta_{eff} \in [0, 0.5]$. This should account both for the Bayes error β_0 of the data that we are given and the additional error $\nu \in [0, 0.5]$ that we introduce by injecting label noise. Then if we wish for the entire β_{eff} portion of the training set to be flagged by q -loss as mislabeled, the empirical risk is $R(\mathbf{w}, b) \geq \beta_{eff} * (1-q)^2$. But we know the trivial solution consisting of all 0 weights has $R(\mathbf{0}, 0) = 1$. Then we want $\beta_{eff} * (1-q)^2 < 1$, which, together with $q \in (-\infty, 0]$, gives $q \in (1 - 1/\sqrt{\beta_{eff}}, 0]$.

Usually we do not have β_0 , but we can obtain an empirical estimate by training on the given data: $\beta_{emp} = \beta_0 + \beta_{opt} + \beta_{gen}$, where β_{opt} is the additional error caused by imperfect optimization, and β_{gen} represents the generalization component of the overall test error. Assuming β_{emp} is sufficiently close to β_0 and accounting for the artificially introduced label noise ν , we set $\beta_{eff} = \beta_{emp} - 2\beta_{emp}\nu + \nu$. The subtraction corrects for originally bad examples that flip under ν .

5. Low-precision discrete variables

The quantum optimization processor that we aim to deploy for solving q -loss requires problems to be discrete and formulated as QUBO. Further, the current hardware can handle a maximum of 512 binary variables, which imposes the additional requirement of being frugal with the bit-depth of weights. To that end we discretize the elements of \mathbf{w} to some low bit-depth $d_w < 64$. While this approach is somewhat unconventional, (Neven et al., 2008) argued there is no fundamental reason why the weights should need high precision and in fact showed a favorable sufficiency condition of $d_w \approx \log(S/N)$ in the case of binary features. Even though classifiers constructed out of more general features have not been studied in this way, our experiments provide support for using low-precision weights.

The reason for fixing at 1 the smallest positive margin that yields zero penalty in q -loss is the same as in hinge loss SVM (Bishop, 2006): any arbitrary rescaling of the weights $\mathbf{w} \rightarrow \kappa\mathbf{w}$ and bias $b \rightarrow \kappa b$ does not change the geometric distance $y_s (\mathbf{w}^T \mathbf{x}_s + b) / \|\mathbf{w}\|$ from a data point (\mathbf{x}_s, y_s) to the decision surface. Therefore, we can assume a margin of 1 for the correctly classified point that is closest to the decision surface. However, this freedom of arbitrary rescaling becomes complicated when the bit-depth of weights is lowered. We want the intervals for weight variables to cover the maximum magnitude that the interplay between margin enforcement and regularization may demand. On the other hand, a loose interval decreases the effective precision in sub-intervals that may really matter. For that reason we derive a λ -dependent bound for setting the intervals in which discrete weight variables can take values.

Let $F(\mathbf{w}, b) = R(\mathbf{w}, b) + \lambda\Omega(\mathbf{w})$ be the objective function. For q -loss, $F(\mathbf{0}, 0) = R(\mathbf{0}, 0) = 1$ and $\exists \hat{\mathbf{w}} \ni F(\mathbf{0}, 0) = \lambda\Omega(\hat{\mathbf{w}})$. Then,

$$F(\hat{\mathbf{w}}, b) = R(\hat{\mathbf{w}}, b) + \lambda\Omega(\hat{\mathbf{w}}) \geq \lambda\Omega(\hat{\mathbf{w}}) = F(\mathbf{0}, 0) \quad (13)$$

Also, $F(-\hat{\mathbf{w}}, b) \geq F(\mathbf{0}, 0)$ because $\Omega(-\hat{\mathbf{w}}) = \Omega(\hat{\mathbf{w}})$. Now consider any $\tilde{\mathbf{w}} \ni \Omega(\tilde{\mathbf{w}}) \geq \Omega(\hat{\mathbf{w}})$:

$$F(\tilde{\mathbf{w}}, b) \geq \lambda\Omega(\tilde{\mathbf{w}}) \geq \lambda\Omega(\hat{\mathbf{w}}) = F(\mathbf{0}, 0) \quad (14)$$

Hence, $F(\tilde{\mathbf{w}}, b) \geq F(\hat{\mathbf{w}}, b) \geq F(\mathbf{0}, 0)$ and $F(-\tilde{\mathbf{w}}, b) \geq F(-\hat{\mathbf{w}}, b) \geq F(\mathbf{0}, 0)$. Thus, we can use $\Omega(\hat{\mathbf{w}}) = 1/\lambda$ to bound the intervals in which the weight variables live while ensuring that the minimizer of F belongs to these intervals. For ℓ_2 -norm regularization, $\Omega(\hat{\mathbf{w}}) = \|\hat{\mathbf{w}}\|_2 \geq \|\hat{\mathbf{w}}\|_\infty = \max(|\hat{w}_j|)$, so we train by optimizing $w_j \forall j$ only in the interval $[-\Omega(\hat{\mathbf{w}}), \Omega(\hat{\mathbf{w}})]$.

The discrete optimization problem for training with q -loss and ℓ_2 -norm regularization is:

$$(\hat{\mathbf{w}}, \hat{b})^* = \arg \min_{\hat{\mathbf{w}}, \hat{b}} \left\{ \frac{1}{S} \sum_{s=1}^S L_q(y_s(\hat{\mathbf{w}}^T \mathbf{x}_s + \hat{b})) + \lambda \|\hat{\mathbf{w}}\|_2 \right\}, \quad (15)$$

where $\hat{\mathbf{w}}$ and \hat{b} are the discretized \mathbf{w} and b , and $\lambda \in \mathbb{R}_{\geq 0}$ controls the relative importance of regularization.

6. Experimental evaluation

While prior work on non-convex losses applied various forms of convex optimization (Masnadi-Shirazi et al., 2010; Yuille & Rangarajan, 2002; Liu et al., 1989) hoping they can still be solved with somewhat reasonable quality, we take the approach of directly tackling the resulting problems by discrete optimization. Admittedly, this choice makes the optimization method largely oblivious to existing benign structure and may cause us to face NP-hardness in certain situations. However, we do this for the purpose of being compatible with emerging quantum hardware that can be employed as a black-box discrete optimization engine having the potential to do well on such problems.

Quantum hardware was already successfully deployed by (Neven et al., 2009a) on a large-scale training problem with square loss and ℓ_0 -norm regularization. In the present work on q -loss with ℓ_2 -norm regularization, we only verify the validity of our approach by using Tabu search (Palubeckis, 2004) as a classical heuristic stand-in and leave the quantum hardware to future work. A quantum optimization with q -loss is expected to achieve in shorter time equal or better results than our classical optimization setup. We do not report CPU time comparisons because they are irrelevant in the absence of quantum hardware runs.

In order to show robustness, we randomly flip training labels and observe the worsening of test error as a function of increasing label noise. While prior work on robust classification (Ding & Vishwanathan, 2010; Collobert et al., 2006; Freund, 2009) considered uniform label noise, we note this does not adequately capture the essence of the true mechanism by which label noise trickles into real-world training tasks. In fact, recent work (Manwani & Sastry, 2011) shows that

even convex losses can be robust under uniform noise. Moreover, experience with practical applications confirms that the type of label noise that affects classification accuracy is never independent of the underlying data distribution. For example, if the human taggers preparing training data for a computer vision application receive somewhat inaccurate or ambiguous instructions affecting only one of the classes, the resulting label noise is strongly correlated with that class. For this reason we move to a noise model in which we introduce uniformly random flips only in the labels of one class—here WLOG of the negative class—and keep the labels of the other class clean. In the experiments described below, the percentage label noise refers to the probability with which we flip labels in the negative portion of training data.

We conduct experiments on two synthetic and four UCI data sets (data summary in Section B of supplementary material). The synthetic data sets (Long & Servedio, 2010; Mease & Wyner, 2007) are designed to provide a stark distinction between robust and non-robust losses. We compare the classification performance of q -loss to seven other convex and non-convex ℓ_2 -regularized methods: liblinear (ℓ_2 -loss primal SVM) (Fan et al., 2008), t-logistic regression (Ding & Vishwanathan, 2010), smoothed hinge loss (Zhang et al., 2010), logistic regression, square loss, sigmoid loss, and probit loss (Bishop, 2006). For all methods except q -loss and liblinear we use Petsc/Tao implementations with convex optimization (Balay et al., 2011; Benson et al., 2010). We do not compare against ramp loss (Collobert et al., 2006), as (Ding & Vishwanathan, 2010) already attempted it in a similar setting on the majority of data sets we use and were unable to produce any salient results. Not surprisingly, this is an example of the inadequacy of convex optimization methods with respect to non-convex problems. Also, we do not compare against 0-1 loss because it is not margin-enforcing. It is well known that if minimized, 0-1 loss yields the lowest possible training error, but due to the lack of margin enforcement, generalization is bad even when regularization is applied (Vapnik, 1998).

With all methods we perform a standard 10-fold cross-validation procedure (Dietterich, 1998) for locating appropriate values of parameters affecting generalization. Fig. 3 presents the main results with an emphasis on the consistently superior performance of q -loss across all data, especially at high levels of noise. We have verified that often in the high noise cases Tabu search fails to reach the lowest attainable objective value. Therefore we believe we are looking precisely at cases of computationally hard optimizations that fail classically but may be solved successfully by quan-

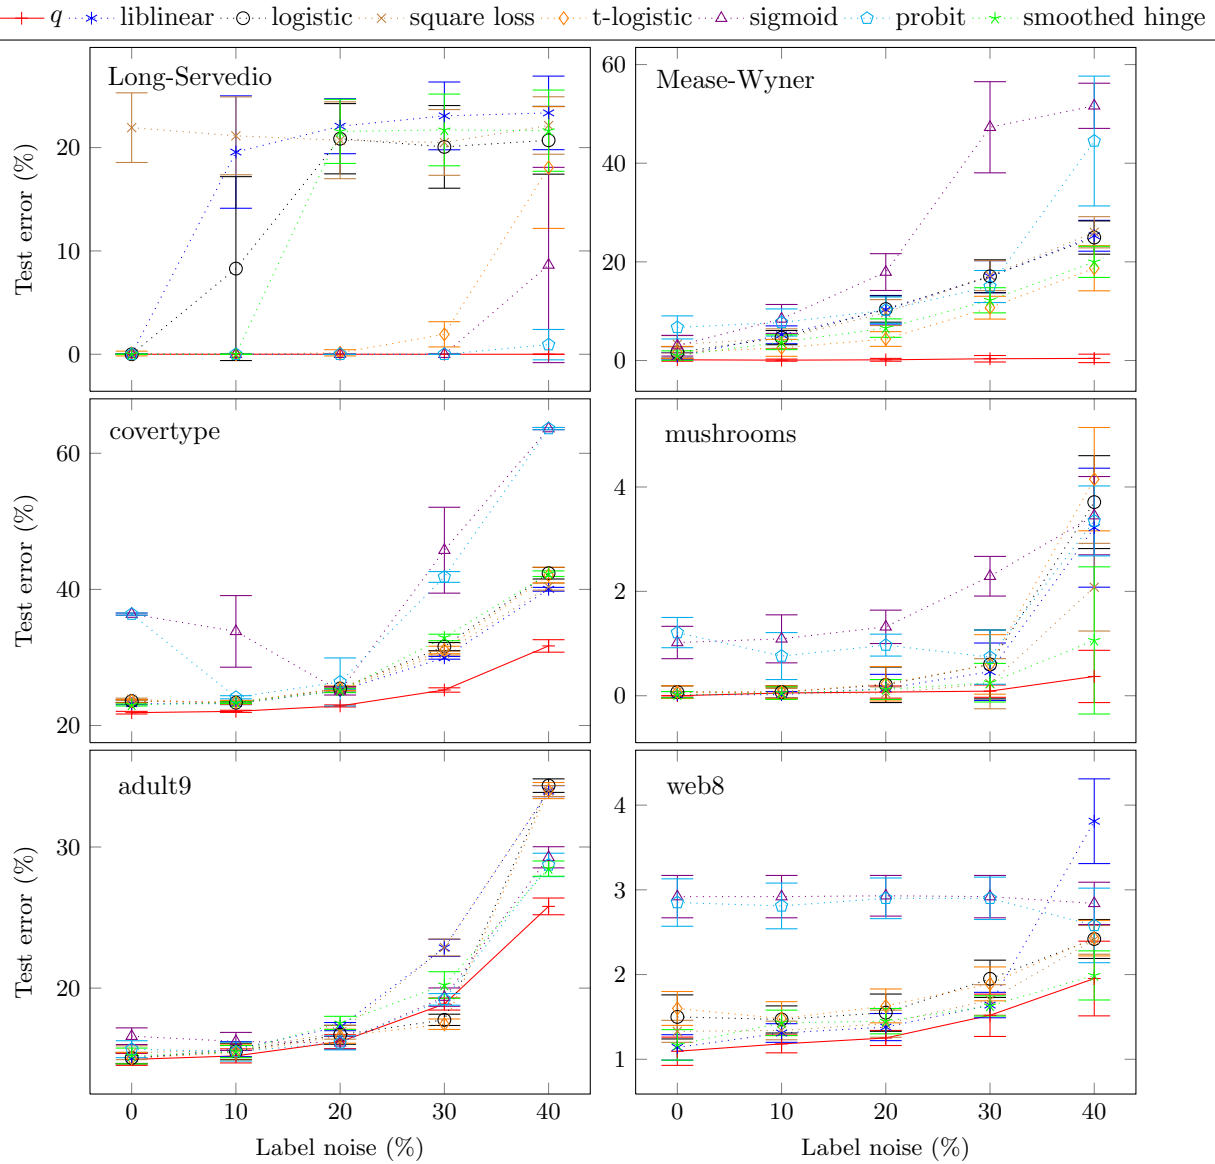


Figure 3. Test error vs label noise for 8 methods (see legend) on 2 synthetic data sets (Long-Servedio and Mease-Wyner) and 4 UCI data sets (covertypes, mushrooms, adult9, web8). Error bars are obtained from 10-fold cross-validation.

tum means. We note sigmoid and probit are sometimes close competitors of q -loss but other times are the worst performers. This can be explained by their non-convexity, which gives them the potential for robustness, but makes them hard to optimize reliably. However, unlike q -loss, we do not know of any AQO-compatible formulations for probit and sigmoid.

q -loss allows us to identify training examples with possibly incorrect labels as the points with $m \leq q$. We recorded the points whose labels we flipped before training (*injected flips*) and the points that q -loss flagged as mislabeled (*trained flips*). Fig. 4 summa-

rizes the overlaps between these two sets. The sets of trained flips for covertypes and adult9 are expectedly larger due to the large Bayes error of these data sets. In the supplementary material we provide details on cross-validated hyper-parameter values (Sections C and D) and statistical significance tests for the observed error rates (Section E).

7. Conclusion

In this paper we introduced q -loss as a robust alternative to convex losses that suffer in the presence of label noise. The QUBO format of the optimiza-




Figure 4. Venn diagrams showing overlap between flips injected in the data before training (injected flips) and flips indicated by training with q -loss (trained flips). Orange color shows portion of injected flips recovered by q -loss.

tion incorporating q -loss makes this version of training an ideal candidate for applying emerging commercial AQO technology as the optimization method of choice. Moreover just by using a classical heuristic solver as a stand-in for hardware-based AQO, we were already able to show significant advantages in test error over a rich variety of data sets and across a number of existing convex and non-convex losses. Our focus here was on formulating a robust loss that can be made compatible with the engineering constraints imposed by emerging quantum hardware. Since with other non-convex losses there is no other choice but to resort to often failing convex optimization, q -loss stands out with its AQO compliance. This opens up new possibilities for achieving results better than ever seen before. Given such encouraging results, we see great potential for robust classification with q -loss under AQO.

Even though (12) is a QUBO, future work still needs to address the fact that on large data sets this formulation may result in a number of binary variables that exceeds the available physical qubits. For that reason, options for training via repeated rounds of optimization—e.g. flavors of large neighborhood search—need to be studied. By using suitable graph embedding techniques, we also need to address the fact that not all quadratic interactions between QUBO variables have corresponding connections between qubits on the physical device. Also, future work needs to investigate the asymptotic scaling of the time necessary for optimizing q -loss with AQO, similarly to the way that was done for square loss in (Neven et al., 2009b). Finally, an interesting open question is whether the derivation in Section 3 can be extended to expressing a more general class of functions as QUBOs.

Acknowledgments

We would like to thank Alessandro Bissacco for providing the “OCR in photos” dataset. We acknowledge Mark Cummins, Edward Farhi, William G. MacReady, James Philbin, and Mani Ranjbar for helpful discussions.

References

- Balay, S., Brown, J., Buschelman, K., Gropp, W. D., Kaushik, D., Knepley, M. G., McInnes, L. C., Smith, B. F., and Zhang, H. PETSc Web page, 2011. <http://www.mcs.anl.gov/petsc>.
- Benson, S., McInnes, L. C., Moré, J., Munson, T., and Sarich, J. TAO user manual (revision 1.10.1). Technical Report ANL/MCS-TM-242, Mathematics and Computer Science Division, Argonne National Laboratory, 2010. <http://www.mcs.anl.gov/tao>.
- Bishop, C. M. *Pattern Recognition and Machine Learning*. Springer-Verlag New York, Inc., 2006.
- Brush, S. History of the Lenz-Ising model. *Reviews of Modern Physics*, 39:883–893, 1967.
- Collobert, R., Weston, J., and Bottou, L. Trading convexity for scalability. In *Proceedings ICML*, 2006.
- Dempster, A. P., Laird, N. M., and Rubin, D. B. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society B*, 39(1):1–38, 1977.
- Dickson, Neil G. and Amin, Mohammad H. Algorithmic approach to adiabatic quantum optimization, 2011. arXiv:1108.3303v1.
- Dietterich, Thomas G. Approximate statistical tests for comparing supervised classification learning algorithms. *Neural Computation*, 10:1895–1923, 1998.
- Ding, N. and Vishwanathan, S. V. N. t-logistic regression. In *Proceedings NIPS*, 2010.
- Ertekin, S., Bottou, L., and Giles, C. L. Ignorance is bliss: Non-convex online support vector machines. *IEEE Transactions on Pattern Recognition and Machine Intelligence*, 33(2):368–381, 2011.
- Fan, R. E., Chang, K. W., Hsieh, C. J., Wang, X. R., and Lin, C. J. LIBLINEAR: A library for large linear classification. *JMLR*, 9:1871–1874, 2008.
- Farhi, E., Goldstone, J., Gutmann, S., and Sipser, M. Quantum computation by adiabatic evolution. 2000. arXiv:quant-ph/0001106v1.

- Feldman, V., Guruswami, V., Raghavendra, P., and Wu, Y. Agnostic learning of monomials by half-spaces is hard. 2010. arXiv:1012.0729v1.
- Freund, Y. A more robust boosting algorithm. Technical report, 2009. arXiv:0905.2138v1.
- Johnson, M. W., Amin, M. H. S., Gildert, S., Lanting, T., Hamze, F., Dickson, N., Harris, R., Berkley, A. J., Johansson, J., Bunyk, P., Chapple, E. M., Enderud, C., Hilton, J. P., Karimi, K., Ladizinsky, E., Ladizinsky, N., Oh, T., Perminov, I., Rich, C., Thom, M. C., Tolkacheva, E., Truncik, C. J. S., Uchaikin, S., Wang, J., Wilson, B., and Rose, G. Quantum annealing with manufactured spins. *Nature*, 473(7346):194–198, May 2011.
- Jordan, Michael I., Ghahramani, Zoubin, Jaakkola, Tommi S., and Saul, Lawrence K. An introduction to variational methods for graphical models. *Machine Learning*, 37:183–233, November 1999.
- Liu, D. C., Nocedal, J., and Dong, C. On the limited memory BFGS method for large scale optimization. *Mathematical Programming*, 45:503–528, 1989.
- Long, P. M. and Servedio, R. A. Random classification noise defeats all convex potential boosters. *Machine Learning Journal*, 783:287–304, 2010.
- Manwani, N. and Sastry, P. S. Noise tolerance under risk minimization. 2011. arXiv:1109.5231v2.
- Masnadi-Shirazi, H., Mahadevan, V., and Vasconcelos, N. On the design of robust classifiers for computer vision. In *CVPR*, pp. 779–786, 2010.
- Mease, D. and Wyner, A. Evidence contrary to the statistical view of boosting. *JMLR*, 9:131–156, 2007.
- Neven, H., Denchev, V. S., Rose, G., and MacReady, W. G. Training a binary classifier with the quantum adiabatic algorithm. 2008. arXiv:0811.0416v1.
- Neven, H., Denchev, V. S., Drew-Brook, M., Zhang, J., Macready, W. G., and Rose, G. NIPS 2009 demonstration: Binary classification using hardware implementation of quantum annealing. December 2009a.
- Neven, H., Denchev, V. S., Rose, G., and MacReady, W. G. Training a large scale classifier with the quantum adiabatic algorithm. 2009b. arXiv:0912.0779v1.
- Palubeckis, G. Multistart tabu search strategies for the unconstrained binary quadratic optimization problem. *Ann. Oper. Res.*, 131:259–282, 2004.
- Vapnik, V. *Statistical learning theory*. Wiley, 1998.
- Yuille, A. L. and Rangarajan, A. The concave-convex procedure (CCCP). In *Proceedings NIPS*, 2002.
- Zhang, X., Saha, A., and Vishwanathan, S. V. N. Regularized risk minimization by Nesterov’s accelerated gradient methods: Algorithmic extensions and empirical studies. 2010. arXiv:1011.0472v1.

A. Explicit QUBO for training problem

Here we explicitly demonstrate the QUBO problem resulting from training with empirical risk over q -loss and Euclidian regularization. To that end we use the variational approximation of q -loss and discretize the optimization variables.

Because the discretization itself is cumbersome and uninstrusive, in Subsection A.1 we first expand all terms with (mostly) continuous variables and show the general layout of the coefficient matrix \mathbf{Q} for the QUBO problem

$$\min_{\boldsymbol{\omega}} \boldsymbol{\omega}^T \mathbf{Q} \boldsymbol{\omega} = \min_{\boldsymbol{\omega}} \left\{ \sum_{\substack{i,j \\ i \neq j}} \omega_i \omega_j [Q_{i,j}] + \sum_i \omega_i [Q_{i,i}] \right\}, \quad (16)$$

where in our case $\boldsymbol{\omega}$ is a concatenation of the binary representations of the discretized N weight variables w , the bias b , and the S variational parameters t . With the notation in (16) we adopt the convention of distinguishing problem coefficients $Q_{*,*}$ by placing them inside square brackets. The preceding symbols are always the corresponding variables.

Finally, in Subsection A.2 we replace the continuous variables \boldsymbol{w} , b , \boldsymbol{t} respectively by their discretized versions $\dot{\boldsymbol{w}}$, \dot{b} , $\dot{\boldsymbol{t}}$ according to bit depths d_w , d_b , d_t and multiplier-offset pairs (α_w, β_w) , (α_b, β_b) , (α_t, β_t) that determine the intervals in which the discrete variables take values.

A.1. Expansion with continuous variables

Using the variational approximation for q -loss, the empirical risk expands as

$$\frac{1}{S} \sum_{s=1}^S L_q(m_s) = \frac{1}{S} \sum_{s=1}^S \min_{t_s} \left\{ m_s^2 - 2m_s t_s + t_s^2 + (1-q)^2 \frac{(1 - \text{sign}(t_s - 1))}{2} \right\}. \quad (17)$$

Now we expand the individual terms appearing on the right-hand side of (17). The goal is to distinguish the coefficients of the various terms in the optimization problem. Hence, in terminal expressions for each term we use the square brackets convention of (16).

$$m_s^2 = (\boldsymbol{w}^T \boldsymbol{x}_s + b)^2 = \sum_{\substack{i=1 \\ j=1}}^N w_i w_j [x_{s,i} x_{s,j}] + bb[1] + b \sum_{i=1}^N w_i [2x_{s,i}] \quad (18)$$

$$-2m_s t_s = -2y_s (\boldsymbol{w}^T \boldsymbol{x}_s + b) t_s = t_s \sum_{i=1}^N w_i [-2y_s x_{s,i}] + b t_s [-2y_s] \quad (19)$$

$$t_s^2 = t_s t_s [1] \quad (20)$$

$$(1-q)^2 \frac{(1 - \text{sign}(t_s - 1))}{2} = (1 - t_{s,d_t})(1-q)^2 = t_{s,d_t} [-(1-q)^2] + (1-q)^2 \quad (21)$$

The idea behind (21) is to use the most significant bit t_{s,d_t} in the binary expansion of t_s as an indicator of $\text{sign}(t_s - 1)$. We give more details on that in Subsection A.2.

The above individual expansions when summed over s become:

$$\sum_{s=1}^S m_s^2 = \sum_{i=1}^N w_i w_j \left[\sum_{s=1}^S x_{s,i} x_{s,j} \right] + bb[S] + b \sum_{i=1}^N w_i \left[2 \sum_{s=1}^S x_{s,i} \right] \quad (22)$$

$$\sum_{s=1}^S -2m_s t_s = \sum_{i=1}^N \sum_{s=1}^S w_i t_s [-2y_s x_{s,i}] + b \sum_{s=1}^S t_s [-2y_s] \quad (23)$$

$$\sum_{s=1}^S t_s^2 = \sum_{s=1}^S t_s t_s [1] \quad (24)$$

$$\sum_{s=1}^S (1-q)^2 \frac{(1 - \text{sign}(t_s - 1))}{2} = \sum_{s=1}^S t_{s,d_t} [-(1-q)^2] + S(1-q)^2 \quad (25)$$

Finally, we can write down the terminally expanded version of the training problem. Note that due to the variational approximation of q -loss, we now have a joint optimization problem over the weight variables \mathbf{w} , the bias b , and the variational parameters \mathbf{t} .

$$(\mathbf{w}, b, \mathbf{t})^* = \arg \min_{\mathbf{w}, b, \mathbf{t}} \left\{ \sum_{i=1}^N w_i w_j \underbrace{\left[\frac{1}{S} \sum_{s=1}^S x_{s,i} x_{s,j} \right]}_{\mathcal{A}_{i,j}} + \underbrace{bb}_{\mathcal{B}} [1] + b \sum_{i=1}^N w_i \underbrace{\left[\frac{2}{S} \sum_{s=1}^S x_{s,i} \right]}_{\mathcal{C}_i} + \sum_{i=1}^N \sum_{s=1}^S w_i t_s \underbrace{\left[\frac{-2y_s x_{s,i}}{S} \right]}_{\mathcal{D}_{i,s}} + b \sum_{s=1}^S t_s \underbrace{\left[\frac{-2y_s}{S} \right]}_{\mathcal{E}_s} + \sum_{s=1}^S t_s t_s \underbrace{\left[\frac{1}{S} \right]}_{\mathcal{F}_s} + \sum_{s=1}^S t_{s,d_t} \underbrace{\left[\frac{-(1-q)^2}{S} \right]}_{\mathcal{G}_s} + \sum_{i=1}^N w_i w_i \underbrace{[\lambda]}_{\mathcal{H}_i} \right\} \quad (26)$$

In (26) we dropped the term $S(1-q)^2$ coming from (25) because it only represents a constant offset. Fig. 5 shows the overall layout of the coefficient matrix implied by the coefficient groups $\mathcal{A}\mathcal{H}$ distinguished from (26). Note that (26) is still using the continuous variables (except for the bits t_{s,d_t}), but it is clear that after discretizing we can obtain the final QUBO from it.

A.2. Binary variables

The final step for obtaining a QUBO is to discretize the continuous variables \mathbf{w} , b , \mathbf{t} via binary expansions of bit-depth d_w , d_b , d_t respectively. We denote the discrete variables by $\dot{\mathbf{w}}$, \dot{b} , $\dot{\mathbf{t}}$. We also define multiplier-offset pairs (α_w, β_w) , (α_b, β_b) , (α_t, β_t) that determine the intervals in which the discrete variables take values.

We apply discretizing transformations by binary variables $w_{*,*}$, b_* , $t_{*,*}$ and the shorthand function $\delta_{\{w,b,t\}}(k) = 2^{k-1} / (2^{d_{\{w,b,t\}}} - 1)$:

$$w_i \rightarrow \dot{w}_i = \alpha_w \sum_{k=1}^{d_w} w_{i,k} \delta_w(k) + \beta_w \text{ for } i = 1, \dots, N \quad (27)$$

$$b \rightarrow \dot{b} = \alpha_b \sum_{k=1}^{d_b} b_k \delta_b(k) + \beta_b \quad (28)$$

$$t_s \rightarrow \dot{t}_s = \alpha_t \sum_{k=1}^{d_t} t_{s,k} \delta_t(k) + \beta_t \text{ for } s = 1, \dots, S \quad (29)$$




Figure 5. Overall layout of the coefficient matrix Q implied by (26) for the QUBO problem $\min_{\omega} \omega^T Q \omega$, where ω consists of the concatenation of the binary representations of the discretized variables \dot{w} , \dot{b} , and \dot{t} . White areas in the matrix correspond to zero coefficients.

The intervals in which the discrete variables take values are:

$$\dot{w}_i \in [\beta_w; \alpha_w + \beta_w]$$

$$\dot{b} \in [\beta_b; \alpha_b + \beta_b]$$

$$\dot{t}_s \in [\beta_t; \alpha_t + \beta_t]$$

As shown in (21) and (25), we take the most significant bit t_{s,d_t} of each variable \dot{t}_s as an indicator for $\text{sign}(\dot{t}_s - 1)$. Therefore we need to choose the interval in which the variables \dot{t} take values such that the most significant bit of each \dot{t}_s is zero for $\dot{t}_s < 1$ and one otherwise. This leads us to intervals for which the upper half of the representable values are greater than or equal to one, so we set $\beta_{\{w,b,t\}} = 1 - \frac{\alpha_{\{w,b,t\}}}{2}$, which gives the intervals in terms of $\alpha_{\{w,b,t\}} > 0$ only:

$$\dot{w}_i \in \left[1 - \frac{\alpha_w}{2}; 1 + \frac{\alpha_w}{2}\right] \quad (30)$$

$$\dot{b} \in \left[1 - \frac{\alpha_b}{2}; 1 + \frac{\alpha_b}{2}\right] \quad (31)$$

$$\dot{t}_s \in \left[1 - \frac{\alpha_t}{2}; 1 + \frac{\alpha_t}{2}\right] \quad (32)$$

Now we convert the various terms in (26) from continuous to binary variables, which gives the coefficients of the

final QUBO problem:

$$\begin{aligned}
 \sum_{i=1}^N \sum_{j=1}^N w_i w_j [\mathcal{A}_{i,j}] &\rightarrow \\
 &\rightarrow \sum_{i=1}^N \sum_{j=1}^N \left(\alpha_w \sum_{k=1}^{d_w} w_{i,k} \delta_w(k) + \beta_w \right) \left(\alpha_w \sum_{k'=1}^{d_w} w_{j,k'} \delta_w(k') + \beta_w \right) (\mathcal{A}_{i,j}) = \\
 &= \sum_{i=1}^N \sum_{j=1}^N \sum_{k=1}^{d_w} \sum_{k'=1}^{d_w} w_{i,k} w_{j,k'} [\alpha_w^2 \delta_w(k) \delta_w(k') \mathcal{A}_{i,j}] + \sum_{i=1}^N \sum_{k=1}^{d_w} w_{i,k} \left[2\alpha_w \beta_w \delta_w(k) \sum_{j=1}^N \mathcal{A}_{i,j} \right] + \beta_w^2 \sum_{i=1}^N \sum_{j=1}^N \mathcal{A}_{i,j}
 \end{aligned}$$

$$\begin{aligned}
 bb[\mathcal{B}] &\rightarrow \\
 &\rightarrow \left(\alpha_b \sum_{k=1}^{d_b} b_k \delta_b(k) + \beta_b \right) \left(\alpha_b \sum_{k'=1}^{d_b} b_{k'} \delta_b(k') + \beta_b \right) (\mathcal{B}) = \\
 &= \sum_{k=1}^{d_b} \sum_{k'=1}^{d_b} b_k b_{k'} [\alpha_b^2 \delta_b(k) \delta_b(k') \mathcal{B}] + \sum_{k=1}^{d_b} b_k [2\alpha_b \beta_b \delta_b(k) \mathcal{B}] + \beta_b^2 \mathcal{B}
 \end{aligned}$$

$$\begin{aligned}
 b \sum_{i=1}^N w_i [\mathcal{C}_i] &\rightarrow \\
 &\rightarrow \left(\alpha_b \sum_{k=1}^{d_b} b_k \delta_b(k) + \beta_b \right) \sum_{i=1}^N \left(\alpha_w \sum_{k'=1}^{d_w} w_{i,k'} \delta_w(k') + \beta_w \right) (\mathcal{C}_i) = \\
 &= \sum_{i=1}^N \sum_{k=1}^{d_b} \sum_{k'=1}^{d_w} b_k w_{i,k'} [\alpha_b \alpha_w \delta_b(k) \delta_w(k') \mathcal{C}_i] + \sum_{i=1}^N \sum_{k'=1}^{d_w} w_{i,k'} [\alpha_w \beta_b \delta_w(k') \mathcal{C}_i] + \\
 &+ \sum_{k=1}^{d_b} b_k \left[\alpha_b \beta_w \delta_b(k) \sum_{i=1}^N \mathcal{C}_i \right] + \beta_b \beta_w \sum_{i=1}^N \mathcal{C}_i
 \end{aligned}$$

$$\begin{aligned}
 \sum_{i=1}^N \sum_{s=1}^S w_i t_s [\mathcal{D}_{i,s}] &\rightarrow \\
 &\rightarrow \sum_{i=1}^N \sum_{s=1}^S \left(\alpha_w \sum_{k=1}^{d_w} w_{i,k} \delta_w(k) + \beta_w \right) \left(\alpha_t \sum_{k'=1}^{d_t} t_{s,k'} \delta_t(k') + \beta_t \right) (\mathcal{D}_{i,s}) = \\
 &= \sum_{i=1}^N \sum_{s=1}^S \sum_{k=1}^{d_w} \sum_{k'=1}^{d_t} w_{i,k} t_{s,k'} [\alpha_w \alpha_t \delta_w(k) \delta_t(k') \mathcal{D}_{i,s}] + \sum_{i=1}^N \sum_{k=1}^{d_w} w_{i,k} \left[\alpha_w \beta_t \delta_w(k) \sum_{s=1}^S \mathcal{D}_{i,s} \right] + \\
 &+ \sum_{s=1}^S \sum_{k'=1}^{d_t} t_{s,k'} \left[\alpha_t \beta_w \delta_t(k') \sum_{i=1}^N \mathcal{D}_{i,s} \right] + \beta_w \beta_t \sum_{i=1}^N \sum_{s=1}^S \mathcal{D}_{i,s}
 \end{aligned}$$

$$\begin{aligned}
 b \sum_{s=1}^S t_s [\mathcal{E}_s] &\rightarrow \\
 &\rightarrow \left(\alpha_b \sum_{k=1}^{d_b} b_k \delta_b(k) + \beta_b \right) \sum_{s=1}^S \left(\alpha_t \sum_{k'=1}^{d_t} t_{s,k'} \delta_t(k') + \beta_t \right) (\mathcal{E}_s) = \\
 &= \sum_{s=1}^S \sum_{k=1}^{d_b} \sum_{k'=1}^{d_t} b_k t_{s,k'} [\alpha_b \alpha_t \delta_b(k) \delta_t(k') \mathcal{E}_s] + \sum_{k=1}^{d_b} b_k \left[\alpha_b \beta_t \delta_b(k) \sum_{s=1}^S \mathcal{E}_s \right] + \\
 &+ \sum_{s=1}^S \sum_{k'=1}^{d_t} t_{s,k'} [\alpha_t \beta_b \delta_t(k') \mathcal{E}_s] + \beta_b \beta_t \sum_{s=1}^S \mathcal{E}_s
 \end{aligned}$$

$$\begin{aligned}
 \sum_{s=1}^S t_s t_s [\mathcal{F}_s] &\rightarrow \\
 &\rightarrow \sum_{s=1}^S \left(\alpha_t \sum_{k=1}^{d_t} t_{s,k} \delta_t(k) + \beta_t \right) \left(\alpha_t \sum_{k'=1}^{d_t} t_{s,k'} \delta_t(k') + \beta_t \right) (\mathcal{F}_s) = \\
 &= \sum_{s=1}^S \sum_{k=1}^{d_t} \sum_{k'=1}^{d_t} t_{s,k} t_{s,k'} [\alpha_t^2 \delta_t(k) \delta_t(k') \mathcal{F}_s] + \sum_{s=1}^S \sum_{k=1}^{d_t} t_{s,k} [2\alpha_t \beta_t \delta_t(k) \mathcal{F}_s] + \beta_t^2 \sum_{s=1}^S \mathcal{F}_s \\
 \sum_{i=1}^N w_i w_i [\mathcal{H}_i] &\rightarrow \\
 &\rightarrow \sum_{i=1}^N \left(\alpha_w \sum_{k=1}^{d_w} w_{i,k} \delta_w(k) + \beta_w \right) \left(\alpha_w \sum_{k'=1}^{d_w} w_{i,k'} \delta_w(k') + \beta_w \right) (\mathcal{H}_i) = \\
 &= \sum_{i=1}^N \sum_{k=1}^{d_w} \sum_{k'=1}^{d_w} w_{i,k} w_{i,k'} [\alpha_w^2 \delta_w(k) \delta_w(k') \mathcal{H}_i] + \sum_{i=1}^N \sum_{k=1}^{d_w} w_{i,k} [2\alpha_w \beta_w \delta_w(k) \mathcal{H}_i] + \beta_w^2 \sum_{i=1}^N \mathcal{H}_i
 \end{aligned}$$

B. Data summary

Table 1. Summary of data sets

Name	Dims	#Examples	Density (%)	Baseline error (%)	d_w
Long-Sevedio	21	2000	100.00	50.00	2
Mease-Wyner	20	2000	100.00	49.80	2
coverttype	54	581012	22.20	36.46	4
mushrooms	112	8124	18.75	48.20	4
adult9	123	48842	11.30	23.93	4
web8	300	59245	4.20	2.92	4

C. q values for q -loss and t values for t -logistic

Table 2. Approximate lower bounds for q in q -loss computed according to Section 5. In each case the q values offered to cross-validation are taken as the 10 equally spaced values between the bound and 0 (both inclusive).

Data set name	Label noise (%)				
	0	10	20	30	40
Long-Servedio	-1000	-3.486401	-2.172365	-1.590225	-1.243201
Mease-Wyner	-25.726124	-3.333979	-2.084934	-1.524450	-1.188680
coverttype	-1.133948	-0.979198	-0.853870	-0.749685	-0.661297
mushrooms	-69.710678	-3.400772	-2.114833	-1.544074	-1.203589
adult9	-1.552383	-1.201211	-0.963926	-0.789849	-0.655127
web8	-8.901475	-2.081794	-1.233931	-0.839671	-0.600122

Table 3. q values for q -loss picked by cross-validation

Data set name	Label noise (%)				
	0	10	20	30	40
Long-Servedio	0	-0.39	-0.24	-0.71	-0.55
Mease-Wyner	0	-2.96	-1.62	-1.36	0
covertime	-0.63	-0.54	-0.38	-0.5	-0.51
mushrooms	0	-0.76	-0.47	-0.17	-0.13
adult9	-0.86	-0.53	-0.43	-0.53	-0.07
web8	-0.99	-0.46	-0.41	-0.19	0

Table 4. t values for t -logistic picked by cross-validation

Data set name	Label noise (%)				
	0	10	20	30	40
Long-Servedio	1.1	1.1	1.5	1.9	2.0
Mease-Wyner	1.1	1.6	2.0	1.9	1.9
covertime	1.1	1.2	1.5	1.3	1.2
mushrooms	1.1	1.1	1.2	1.1	1.1
adult9	1.2	1.3	2.0	1.2	2.0
web8	1.1	1.1	1.1	1.2	1.1

D. Regularization strength

Table 5. λ values offered to cross-validation (C for liblinear is $1/\lambda$)

λ
2.000090
0.398965
0.079583
0.015875
0.003167
0.000632
0.000126
0.000025
0.000005
0.000001

Table 6. C values for liblinear picked by cross-validation

Data set name	Label noise (%)				
	0	10	20	30	40
Long-Servedio	0.499978	2.506486	0.499978	0.499978	0.499978
Mease-Wyner	40000.000000	0.499978	315.756236	12.565498	62.992126
covertime	0.499978	2.506486	62.992126	1000000.000000	12.565498
mushrooms	2.506486	12.565498	0.499978	0.499978	0.499978
adult9	0.499978	62.992126	0.499978	0.499978	0.499978
web8	315.756236	0.499978	12.565498	12.565498	0.499978

Table 7. λ values picked by cross-validation for 0% label noise

Data set name	Method						
	q	logistic	square	t -logistic	sigmoid	probit	smooth hinge
Long-Servedio	0.015875	0.003167	0.079583	0.003167	0.000632	0.003167	0.015875
Mease-Wyner	0.000126	0.000001	0.000025	0.000001	0.000025	0.003167	0.000001
covertime	0.000025	0.000025	0.000025	0.000001	2.000090	2.000090	0.000001
mushrooms	0.000025	0.000001	0.000025	0.000126	0.000632	0.015875	0.000632
adult9	0.003167	0.000001	0.000632	0.000025	0.003167	0.003167	0.000126
web8	0.000632	0.000001	0.000005	0.000005	0.000126	0.000632	0.000001

Table 8. λ values picked by cross-validation for 10% label noise

Data set name	Method						
	q	logistic	square	t-logistic	sigmoid	probit	smooth hinge
Long-Servedio	0.015875	0.000005	2.000090	0.000126	0.000632	0.003167	0.003167
Mease-Wyner	0.000126	0.000005	0.000632	0.000001	0.000632	0.003167	0.000005
covertypes	0.000001	0.000025	0.000632	0.000001	0.000632	0.003167	0.000126
mushrooms	0.003167	0.000005	0.000001	0.000001	0.000632	0.003167	0.000005
adult9	0.015875	0.000632	0.003167	0.000632	0.003167	0.003167	0.000126
web8	0.000632	0.000005	0.000126	0.000001	0.000126	0.000632	0.000005

Table 9. λ values picked by cross-validation for 20% label noise

Data set name	Method						
	q	logistic	square	t-logistic	sigmoid	probit	smooth hinge
Long-Servedio	0.000126	2.000090	2.000090	0.000025	0.000632	0.003167	2.000090
Mease-Wyner	0.000126	0.000025	0.000005	0.000001	0.003167	0.003167	0.000126
covertypes	0.000001	0.000001	0.000126	0.000001	0.000025	0.000632	0.000126
mushrooms	0.003167	0.000126	0.000632	0.000126	0.000632	0.003167	0.000025
adult9	0.015875	0.079583	0.079583	0.003167	0.000632	0.003167	0.003167
web8	0.000632	0.000001	0.000001	0.000005	2.000090	0.000632	0.000126

Table 10. λ values picked by cross-validation for 30% label noise

Data set name	Method						
	q	logistic	square	t-logistic	sigmoid	probit	smooth hinge
Long-Servedio	0.003167	2.000090	2.000090	0.000001	0.000126	0.003167	2.000090
Mease-Wyner	0.000126	0.000005	0.000001	0.000001	0.003167	0.003167	0.000005
covertypes	0.000025	0.000001	0.000126	0.000001	0.000632	0.003167	0.000025
mushrooms	0.003167	0.000632	0.003167	0.000632	0.003167	0.003167	0.000632
adult9	0.003167	2.000090	0.003167	2.000090	0.000126	0.000632	2.000090
web8	0.000632	0.000126	0.000001	0.000632	0.000632	0.003167	0.000126

Table 11. λ values picked by cross-validation for 40% label noise

Data set name	Method						
	q	logistic	square	t-logistic	sigmoid	probit	smooth hinge
Long-Servedio	0.003167	2.000090	2.000090	0.000001	0.000126	0.000632	2.000090
Mease-Wyner	0.000126	0.000001	0.000005	0.000001	2.000090	0.003167	0.000001
covertypes	0.000001	0.000001	0.000001	0.000001	2.000090	2.000090	0.000001
mushrooms	0.003167	0.000126	0.000632	0.000001	0.003167	0.015875	0.003167
adult9	0.000025	0.000126	0.000126	0.000001	0.000632	0.003167	0.079583
web8	0.000632	0.015875	0.079583	0.015875	0.000632	0.000632	0.000632

E. Statistical significance

Table 12. Paired t -test for statistical significance of the test error difference yielded by q -loss and an estimated closest competitor. The closest competitor is manually chosen on a per-data-set basis from the set of convex losses. We exclude the other non-convex losses (t-logistic, sigmoid, and probit) from this comparison as we do not know what their performance would be if they could be reliably solved to optimality. We reject the null hypothesis at $\alpha = 5\%$ significance level. 'Y' means that the difference is significant and 'N' means the difference is not significant.

Data set name	Compared losses	Label noise (%)				
		0	10	20	30	40
Long-Servedio	smoothed hinge vs q -loss	N	N	Y	Y	Y
Mease-Wyner	smoothed hinge vs q -loss	Y	Y	Y	Y	Y
covertypes	liblinear vs q -loss	Y	Y	Y	Y	Y
mushrooms	smoothed hinge vs q -loss	N	N	N	N	Y
adult9	smoothed hinge vs q -loss	Y	Y	Y	Y	Y
web8	smoothed hinge vs q -loss	Y	Y	Y	Y	Y